NASA-TM-112079

7N-62-TM
016257

[To be presented at the Third International Conference on Supercomputing, Boston Massachusetts, May 15-20, 1988]

STATUS AND FUTURE DEVELOPMENTS OF THE NAS PROCESSING SYSTEM NETWORK

RND-88-001

Bruce T. Blaylock

F. Ron Bailey.

Numerical Aerodynamic Simulation Systems Division

NASA Ames Research Center, Moffett Field, California 94035

1.0 Abstract:

NASA has an extensive and on-going program in computational aerodynamics and fluid dynamics. The Numerical Aerodynamic Simulation Program was established at the NASA Ames Research Center to provide a computational capability and facility in support of a nationwide user base. The Initial Operating Configuration was achieved in March 1987 and included a Cray 2, Silicon Graphics scientific workstations and several support processors interconnected both locally and remotely using high performance network technologies. The Extended Operating Configuration planned for July 1989 will provide an overall increase in system throughput from the current 250 MFLOPS to more than 1 GFLOPS. The Extended Operating Configuration will include addition and improved capabilities over those currently provided.

2.0 Introduction

The NAS Program origins date from 1975, when ARC researchers in Computational Fluid Dynamics (CFD) began to investigate the required computer technologies needed to support the advances then being made in fluid dynamics modeling. For a historical overview of the NAS Program see ref. 1 and for overall Program status see ref. 2. After nine years of program advocacy and several technical studies, the NAS Program became a NASA New Start Program in Fiscal Year 1984 with the following objectives:

- o To provide a national computational capability, available to NASA, DoD, other government agencies, industry, and universities to ensure the Nation's continuing leadership in computational fluid dynamics and related disciplines.
- o To act as a pathfinder in advanced large scale computer system capability through systematic integration of state-of-the-art improvements in computer hardware and software technologies.
- o To provide a strong research tool for the NASA Office of Aeronautics and Space Technology

3.0 NAS Processing System Network Description (NPSN)

The original NAS Program Plan, approved in February 1983, called for the development of computational capability able to provide a sustained computational rate of 1 GFLOPS (billion floating-point operations per sec). Given the estimation of hardware performance and availability, the NAS Program was split into two development phases. These two phases culminated in the Initial Operating Configuration (IOC) to be achieved in 1986 with 250 MFLOPS (million floating-point operations per sec), and the Extended Operating Configuration (EOC) to be achieved in 1987 with 1 GLOPS. The original IOC system diagram is shown in Figure 1. The diagram illustrates the following components:

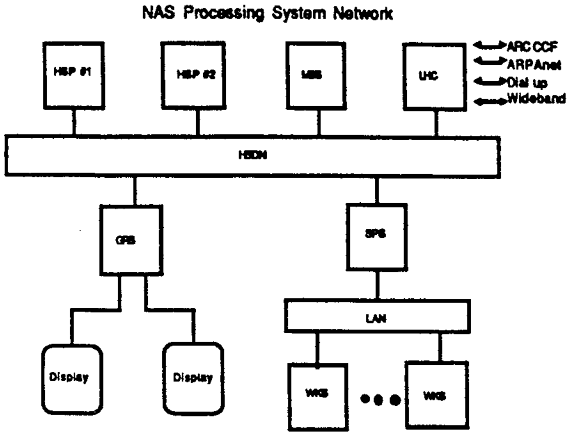

Figure 1 Original NPSN System Design

1. High Speed Processors (HSP) targeted to be large-scale, ultra-high speed, scientific computer systems which were to provide the computational throughput and large memory capacity required to process ever more advanced numerical flow simulation programs. The first high speed processor was to be capable of providing peak computational rates of two billion floating point operations per second, a sustained rate of 250 million floating point operations per second, and a central memory of at least 64 million 64-bit words (expandable to 256 million 64-bit words). The figure also illustrates the NAS strategy of always having two high speed processors, with the oldest, most stable high speed processor in production use and the newest one in the development stage.

2. A Support Processing Subsystem (SPS) targeted to provide general purpose processing for program development, result editing and other ancillary support applications. In addition the SPS was to provide a network gateway for the workstation subsystem.

3. A Mass Storage Subsystem (MSS) to provide on-line and archival storage, and to function as the main user data storage and file management subsystem.

4. A Graphics Subsystem (GRS) to provide local users with a system combining local processing and storage with sophisticated, dynamic graphics display devices .

5. A Workstation Subsystem (WKS) defined as a collection of intelligent terminals, local processors and associated storage devices. The WKS was to provide interactive terminal functions, local processing and limited graphics displays.

6. A High Speed Data Network Subsystem (HSDN) to provide complete interconnection between the HSP, SPS, GRS, and MSS, primarily in support of large file transfers.

7. A Local Area Network subsystem (LAN) targeted to be the physical data transfer path between the SPS and the WKS, providing interactive terminal and file transfer communication.

8. A Long Haul Communication Subsystem (LHC) serving to connect remote users of the NAS with the NPSN. The types of connectivity anticipated included access to the NPSN from the ARC central computing facility, the ARPAnet, dial-up service, and wide band communications.

As the design and development of the NAS Processing System Network evolved, the target (IOC) design was modified as shown in Figure 2. The figure illustrates the collection of the SPS, MSS and LHC functions into a single entity called the Integrated Support Processing Complex (ISPC). The second major change shown in Figure 2 is the direct connection of the workstations to the high speed data network with the elimination of the Support Processing Subsystem as a gateway.

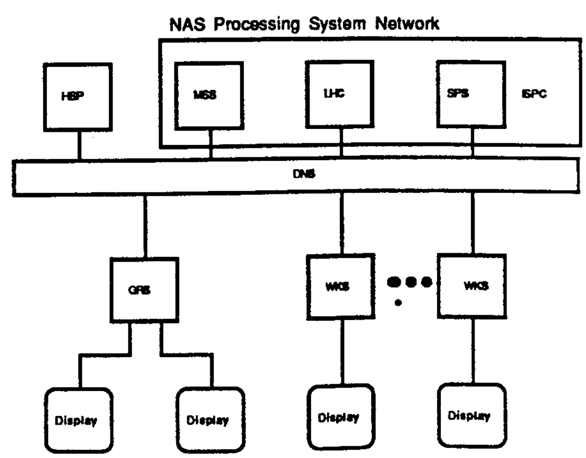

Figure 2 Target IOC NPSN System Design

The NAS Processing System Network design and development continued to mature as systems were installed and integration took place. For example, the GRS was folded into the increasingly powerful workstation subsystem.

On March 9, 1987, the complete capability provided by the NAS Program was formally dedicated and the functional Initial Operating Configuration was declared . Figure 3 illustrates the system configuration at the time of IOC.

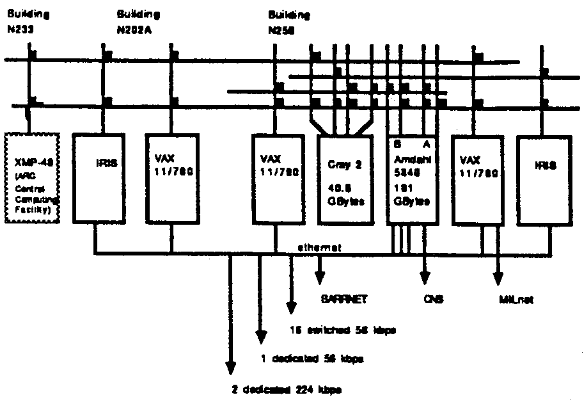

Figure 3 Final IOC NPSN System Design

The NAS system at IOC contained the following major hardware and software elements:

1. A Cray 2 serving as the High Speed Processor. The Cray 2 has four processors, 268 million words (decimal) of central memory, 34 DD49 disk drives, and four Network Systems Corporation (NSC) A-130 HYPERchannel adapters. The Cray 2 runs the UNICOS operating system, a derivative of AT&T System V UNIX with additional networking software supplied by Cray Research, Inc. (CRI).

2. An Integrated Support Processor Complex providing the mass storage and support processing functions on two Amdahl 5840 processors configured with 190 GBytes of DASD, five NSC A-222 HYPERchannel adapters and three Spartacus Ethernet adapters. The Amdahl processors run the VM Operating System with multiple guest virtual operating systems under VM. The user-visible operating system is UTS/V, the Amdahl UNIX System V (version 2) product, together with networking software supplied by the Wollongong Group, Inc.

3. 35 Silicon Graphics 2500 Turbo and 3030 IRIS Workstations. The operating system on the workstations is a NAS-modified version of the Silicon Graphics GL2.3 operating system and includes Berkeley-UNIX networking features.

4. Four VAX 11/780 computers. Of these, one is for general production use and acts as the primary Milnet gateway. Another is used for development and serves as a secondary Milnet gateway. The third is dedicated to production use by scientific users. The fourth VAX is used as a user-transparent long haul network gateway, for limited development, and as the source machine for the other VAXes. The operating system on the VAX computers is 4.3 bsd UNIX.

5. Two Local Area Network Technologies, ethernet and Network Systems Corporation HYPERchannel, provide the physical interconnection mediums for the system. With the exception of the Cray computers all the processors are connected with both technologies. The networks span three facilities at Ames and cover a distance in excess of a mile. The HYPERchannel network has been installed as four subnets with each subnet using two coaxial trunks.

6. The communication with remote user sites is accomplished using switched and dedicated high speed data lines coupled to Vitalink Corporation TransLan devices. The Vitalink equipment allows for the interconnection of ethernets to and creates the logical equivalent of a nationwide "local area network". As of December 1987, over 700 researchers at 75 remote sites were directly tied into the NAS System via remote communications links. In addition, interconnection to other computer networks including BARRnet and MILnet has been achieved. A bulk file transfer capability between selected NASA centers called CNS (Computer Network System) is also supported.

4.0 Extended Operating Configuration

As with IOC, the major feature that defines the NAS EOC is the high speed processor that will serve as the focal point for the rest of the NPSN system. The actual date for establishing EOC is dependent on the delivery, installation, testing and acceptance schedules for the new high speed processor, known as HSP 2. However, during both the design and development phases of the IOC NPSN and during its initial operation, several other necessary and important improvements and additions to the NPSN have been identified. Although perhaps not as dramatic as the addition of a second high speed processor, these new features will also make a significant impact on the level of service to be provided to the scientific users of the NAS EOC system.

Improvements and upgrades to existing NPSN subsystems will include extensive improvements to the Mass Storage Subsystem (MSS), the development and incorporation of a high speed local area network (High Speed LAN), reconfiguration of the NPSN Ethernet network to isolate remote traffic from local traffic, and the introduction of new, more powerful scientific workstations. Several major new additions will also be incorporated into the EOC NPSN, including

Mini-Supercomputers and an Auxiliary Processing Center (APC).

The first important change in the overall NPSN configuration for EOC is the introduction of multiple mini-supercomputers. Although the mini-supercomputers have been designated to perform a wide variety of functions for EOC, the original impetus for including them in the EOC design was a re-examination of the Computational Fluid Dynamics (CFD) tasks and how these tasks are mapped onto the existing NPSN processors.

In a somewhat simplified view of CFD work, there are three major driving tasks:

1. Pre-processing in preparation for solution generation (which includes grid development, body definition, specification of boundary conditions, algorithm development, and symbolic computation). Pre-processing activities are limited to a time scale imposed by human interaction and tend to be scalar in nature.

2. Solution generation in terms of the computation of some form of the Navier Stokes equations and the production of a solution set. These activities include steady Reynolds-averaged Navier Stokes, large eddy simulation, and unsteady Reynolds-averaged Navier Stokes. Solution generation activities are primarily CPU-intensive vector computations that require little if any human intervention.

3. Post-processing analysis of generated solutions. Postprocessing activities include graphics rendering, data base interrogation (as in particle path calculation),and fast digital playback of fluid flow animation sequences. Post-processing analysis activities are again limited to a time scale imposed by human interaction and tend to be scalar in nature.

The mapping of these activities onto the IOC NPSN system is shown in Figure 4; the major portions of the current system shown in the diagram include the high speed processor (Cray 2) and the workstations (IRIS 2500 Turbo and 3030s).

As the figure shows, in the current working model some portion of all three major CFD tasks is now performed on the Cray 2. This multiplicity of activities on the Cray 2 is the direct result of its processing power when compared to the existing IRIS workstations, which have been used primarily for some pre-processing activities and as the flow visualization tool. However, this distribution of function requires the Cray 2 to be both a scalar and a vector processor, and to support both CPU-intensive and interactive tasks.

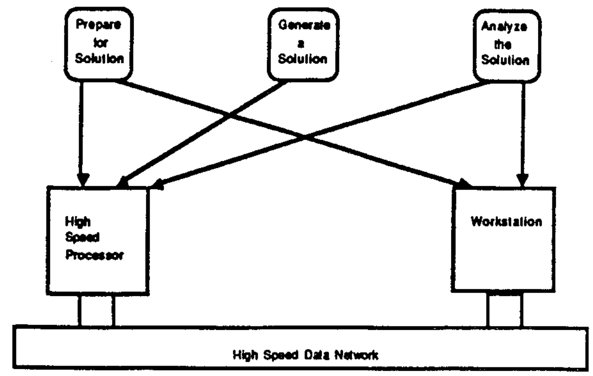

Figure 4 Current Mapping of the CFD Tasks Onto the NAS Processors

Both high speed processors at EOC (the Cray 2 and HSP 2) will collectively represent the most expensive and the least easily replicated resource of the NPSN, and the goal for EOC is to redistribute the processing load between the high speed processors and other NPSN machines. Redistribution of the scalar and interactive CFD tasks to mini-supercomputers with performance capabilities somewhere between those of the high speed processors and the workstations will make most efficient use of the NPSN high speed processing resources.

Figure 5 shows how the addition of mini-supercomputers allows for a repartitioning of the CFD workload. The interactive and scalar activities have been shifted to the mini-supercomputers, leaving the high speed processor(s) to perform CPU-intensive vector processing. Further it is anticipated that the working storage requirements will be best served by either the mini-supercomputers or the workstations with the Mass Storage Subsystem relegated to archival storage.

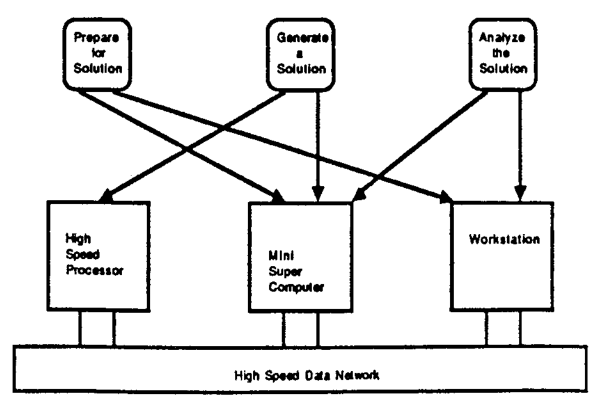

Figure 5 Proposed Mapping of the CFD Tasks Onto the NAS Processors

The logical reconfiguration of the NPSN for EOC is illustrated in Figure 6 The figure shows clusters of N workstations associated with a smaller number (M) of mini-supercomputers; both workstations and mini-supercomputers are interconnected by multiple bandwidth LAN technologies to the high speed processors. It is assumed that the number of high speed processors (L) will always be 1 or 2, the number of workstations per cluster (N) will range between 1 and 10, and the number of mini-supercomputers (M) will range from 5 to 20 depending on 1) the specific functionality of each mini-supercomputer, 2) the total number of workstations being supported, and 3) the number of workstations in the clusters.

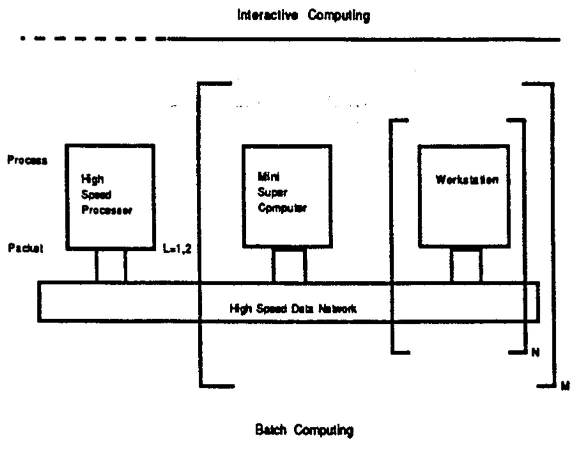

Figure 6 Logical NPSN Design for EOC

Although the specific functions to be performed by each of the 'M' mini-supercomputers to be incorporated into the EOC NAS system has not yet been determined, it is expected that several will be assigned to research scientists as 'scientific' machines and that others will perform the functions currently assigned to the Support Processing Subsystem. As described below, several will also form an important part of the configuration in support of secure processing.

4.1 EOC NPSN Hardware

A functional NPSN configuration for EOC is provided in Figure 7. The diagram shows local and remote workstation clusters, each with an associated mini-supercomputer. Two clusters of workstations and associated mini-supercomputers are shown in the Auxiliary Processing Center (APC).

4.2 EOC NPSN Software

Software that resides on the various processors within the NPSN can be classified into one of three general categories:

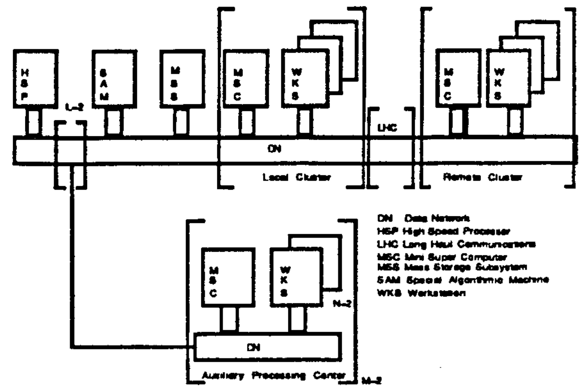

Figure 7 EOC NPSN Target Functional Configuration

1. COMMON software, such as utilities and other programs that have the same form and function on several processors NPSN COMMON software includes such items as full screen editors (e.g. EMACS), command languages (e.g. 'C' Shell), and various compilers

2. DISTRIBUTED software, or software that utilizes the network to achieve interprocess communication The NPSN DISTRIBUTED software for EOC includes such features as the 4.3 bsd UNIX 'r' commands (rlogin, rcp, rwho, etc.), flow visualization tools (PLOT3D, GAS), and electronic mail. A bulk data transfer protocol will be developed for EOC, and specialized windowing software, such as NeWS or Xwindows, will be selected and implemented.

3. SUBSYSTEM-SPECIFIC software, or software that executes on only one processor in the network The NPSN SUBSYSTEM-SPECIFIC software includes the operating systems (all UNIX, but different varieties), special scientific libraries on the high speed processors, and network drivers.

One of the major software additions for EOC will be the acquisition of Network File System (NFS) as part of the NPSN DISTRIBUTED software. NFS is a product of Sun Microsystems and is a widely implemented alternative to the AT&T remote file system. NFS provides transparent file access among computers from many different manufacturers over one or more networks. Files can reside on high-performance workstations, minicomputers, mainframes and even personal computers and be commonly shared. Different file structures and operating systems are transparent.

A major software acquisition will be required in order to upgrade the existing UNIX operating systems on the NPSN processors. The UNIX operating system was originally designed and developed within Bell Laboratories and remains a licensed product of AT&T. The right to install UNIX, either in source or binary form, on any NPSN computer is dependent on the existing licensing arrangements held by the vendors and NASA. The current UNIX license base for all NPSN computers is the AT&T System V release 2 version for the VAX 11/780. In addition, site licenses exist for both 4.2 bsd and 4.3 bsd UNIX.

4.3 Development Activities Leading to the Extended Operating Configuration

The development of the Extended Operating Configuration will be accomplished through the addition of new capabilities and the incremental improvement to existing capabilities. These developments are described in the following sections.

4.3.1 High Speed Processor 1 (Cray 2) Development

For EOC the major HSP development effort will be directed toward introducing the second high speed processor (HSP 2) into the NPSN configuration. However, some additional effort will be directed toward expanding and upgrading the capabilities provided by HSP 1 (the Cray 2). For EOC, the Cray 2 will provide between 20 and 25% of the NPSN High Speed Processor computing resources. In addition, the Cray 2 will provide a development platform for prototype and pathfinding technology, such as high speed LANs. It will also be the High Speed Processor resource for the Auxiliary Processing Center.

The major Cray 2 development effort for EOC will focus on providing a higher level of service to NPSN users; high user availability and reliability will be key objectives. Performance improvements will center on new products introduced by Cray Research, Inc. (CRI) and other vendors, and on software tuning and system management. The level of effort devoted to system tuning in particular will rise as usage of the Cray 2 increases and the introduction of new products decreases.

The anticipated configuration for the Cray 2 serving as High Speed Processor 1 at the time of EOC could include:

- o 4 background CPUs
- o 256 MW of 80 nanosecond dynamic RAM memory
- o 12 DD40 disks
- o 21 DD49 disks
- o 3 N-Series HYPERchannel adapters (one for the dedicated MSS-Cray 2 connection, one for connection to the workstations, and one for mini-supercomputer and SPS connections. The third adapter will also serve as a backup with reconfiguration).
- o 1 LAN interface supporting from two to four Ethernets and one Pronet
- o 1 Cray tape channel with two IBM controllers and four 3480 tape drives
- o 1 High Speed External channel (HSX)
- o 1 channel which would be used either for a second HSX connection or a second LAN connection if CRI provides support for only two Ethernets per Ethernet interface

The configuration for High Speed Processor 1 as of January 1988 is shown in Figure 8.

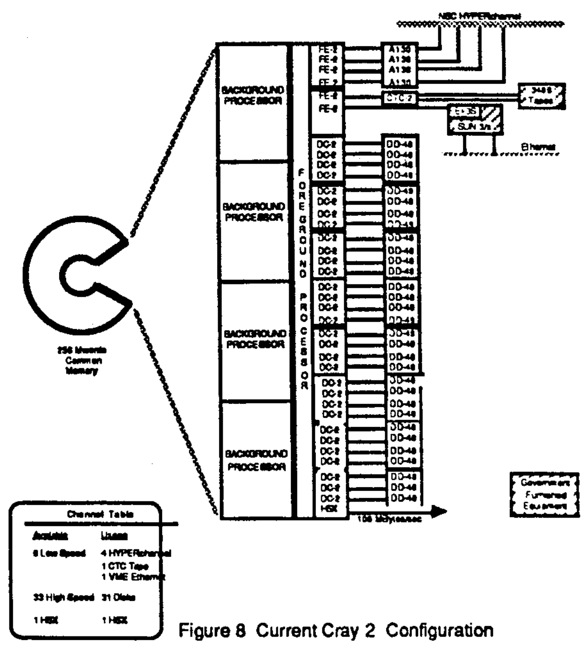

Figure 8 Current Cray 2 Configuration

EOC software development activities for the Cray 2 will combine local software development and support of existing software with software provided and supported by CRI. The CRI software will consist of scheduled releases of UNICOS, with releases expected to occur every six months. In some cases, major functions introduced in these releases of UNICOS may require several months of local development effort before the software can be placed in production use.

For UNICOS releases 3.0, 4.0, and 5.0, the major features to be incorporated are already known and are outlined below. The features to be included in later releases are as yet undefined.

UNICOS 3.0 , incorporated into the Cray 2 software configuration in July, 1987, provided job checkpoint and restart capability for Network Queuing System (NQS) jobs, UNIX standard performance accounting, Berkeley 4.3 style subnet routing, improved multitasking support, Xwindows, and tape support for 3480 cartridge tape drives. In October 1987, several language upgrades were incorporated, including better local memory for CFT, vectorization and scalar optimization for C, and multitasking capability for Pascal. In December 1987, CFT-77 improvements included automatic partitioning for multitasking, improved vectorization, the INCLUDE statement, and various performance improvements.

UNICOS 4.0 may be available in the January/February 1988 timeframe. This release will incorporate improved tape support including a tape daemon for resource allocation, microtasking, HSX support, and Network File System (NFS) with its underlying layers of Remote Procedure Calls (RPC) and External Data Interchange (XDR). It is anticipated that NFS will become operational several months after 4.0 is put into production, since it will require testing and trial use before being made available to users. Release 4.0 of UNICOS will also provide additional capabilities currently provided by COS but as yet not available in UNICOS.

UNICOS 5.0 should be available in July 1988. It will provide tape archiving capability that is similar in function to COS tape archiving, and support for the N-series HYPERchannel adapters (although CRI has made no firm commitment to this).

In addition to releases of UNICOS, EOC software development effort for the Cray 2 will focus on several areas needing improvement. An operational backup utility that reliably and efficiently moves data from disk and allows its restoration will be a critical need in the EOC timeframe. Currently, the Integrated Backup and Recovery (IBR) software is being used to perform this function on the SPS, and will be assessed for its applicability in providing this function on the Cray 2.

Accounting software on the Cray 2 is also another area targeted for improvement by EOC. CRI has made a number of changes to the user account data format, but to date has not announced any specific plans to improve the existing accounting software. As a result, it may be necessary to continue local improvements for usage accounting and reporting.

Work in the areas of creating and improving general performance analysis tools will be on-going. In particular, trend analysis and system throughput analysis will become increasingly important as more users and an increasing number of jobs are introduced into the system.

Additional user tools will be developed so that the programming process will become more efficient for the scientific users of the Cray 2. This software will include program analysis tools for optimizing and vectorizing user programs. Eventually, user tools will exist to provide automatic multitasking, either with pre-processor or with compiler support. It is also possible that NAS resources will be required to improve and enhance multitasking support for UNICOS at the operating system and program support levels.

Other anticipated software development tasks for EOC include queueing support for hardcopy routing, improved operator control, and new graphics utilities.

4.3.2 High Speed Processor 2 Development

The major element of the NAS Extended Operating Configuration is the addition of a second High Speed Processor (HSP) to the NPSN. A summary of the function and performance requirements for the HSP 2 system is contained in Table 1.

Table 1. High Speed Processor Final System Requirements

Hardware

Tested Peak Speed	2	Gflops
Sustained Speed	1	Gflops
Memory Capacity	2	GBytes 256MW
Disk Capacity	50	GBytes
Peak Speed / Processor	200	Mflops
Disk to Memory		
(/dev)	8	MBytes / sec
(total)	200	MBytes / sec
Sustained Speed on Benchmark		
per processor	100	Mflops
Total	1	Gflops
# HYPERchannel interfaces	6	
# Ethernet interfaces	6	
Adapter to adapter rates		
HYPERchannel	5	MBytes / sec
Ethernet	1	MBytes / sec
High Speed Interface	100	MBytes / sec

Software

UNIX System V release 2, NQS, Berkeley Networking, Fortran

Since the performance requirements for the second high speed processor are at the leading edge of current technology, it is anticipated that the HSP acquired will be one of the first of its kind. Because low serial number machines present special problems, a risk reduction strategy has been adopted in which the HSP 2 will be acquired through a two step process.

The competitive procurement strategy allows for the possible award of multiple contracts. The selected vendors will install initial systems of a reduced capability in the NPSN for the purpose of initial evaluation. Based on the performance of the initial systems a determination will be made as to which of the initial systems will be upgraded to a final HSP 2 configuration. The performance requirements called for in an initial system are outlined in Table 2.

Table 2. High Speed Processor Initial System Requirements

Hardware		
Tested Peak Speed	500	Mflops
Sustaine d Speed	250	Mflops
Memory Capacity	500	MBytes
	64 MW	
Disk Capacity	25	GBytes
Peak Speed / Processor	100	Mflops
Disk to Memory		
(/dev)	8	MBytes / sec
(total)	200	MBytes / sec
Sustained Speed on Benchmark		
per processor	50	Mflops
Total	250	Mflops
# HYPERchannel interfaces	6	
# Ethernet interfaces 6		
Adapter to adapter rates		
HYPERchannel	5	MBytes / sec
Ethernet	1	MBytes / sec
High Speed Interface	100	MBytes / sec
Software		
UNIX System V release 2, NQS, Berkeley Networking., Fortran		

As of December 1987, one HSP 2 contract had been awarded to Cray Research Inc and second contract with ETA Systems was still being negotiated.

4.3.3 Mass Storage Development

The Mass Storage Subsystem (MSS) hardware consists of one of two Amdahl 5840 processors, 120 GBytes of DASD, IBM 3480 cartridge tape drives, HYPERchannel and Ethernet network connections, and an IBM 3088 channel-to-channel adapter.

The Amdahl 5840 processor used for the MSS contains 24 MBytes of memory and 24 channels. The CPU configuration of the MSS Amdahl and its channel connections are identical to those of a companion Amdahl 5840, which is used primarily to provide the Support Processing Subsystem (SPS) function. This design enables either Amdahl to provide the capabilities of both the SPS and the MSS in the event of a catastrophic failure of one of the two CPUs.

The current MSS has access to 120 GBytes of rotating mass storage, composed of six strings of Amdahl 3480-compatible single density DASD, with two independent channel paths to each string.

The second level of MSS storage is provided by eight manually operated IBM 3480 cartridge tape drives, connected through two controllers and two channels. Cartridge tapes are manually stored and retrieved from a collection of tape storage racks capable of holding approximately 60,000 tapes.

Network connections are provided by two HYPERchannel adapters, each connected to its own channel, and a Spartacus Ethernet adapter that shares access to a channel with 3270 terminal controllers.

The current MSS software is based on a close relationship between the UTS and MVS guest operating systems, both operating under VM on the same Amdahl CPU. The two guest operating systems share the functions necessary to fulfill the requirements for an MSS, with UTS providing the user interface through network connections, and MVS managing the data storage.

The software running under UTS consists of TCP/IP software from The Wollongong Group Inc., a locally developed HYPERchannel driver, a pseudo-disk driver that makes the channel-to-channel connection to MVS look like disk space to UTS, and various utilities and changes to utilities to account for the other additions to UTS. In addition, a virtual HYPERchannel adapter package was written for VM, so multiple guest hosts can share access to a single physical HYPERchannel adapter.

The software running under the MVS operating system consists of a "parent"-initiated task establishing communication with UTS, its "children", and the DMS/OS hierarchical storage manager. The "parent"-initiated task communicates with UTS through the channel-to-channel interface (either real or virtual), and farms out "child" tasks to carry on I/O between UTS and individual files.

MSS files are mapped into one or more fragmented VSAM data sets. A special "child" task examines closed files for efficient organization and moves them between the large fragment pool and the small fragment pool. Finally, the "parent" task communicates with DMS/OS (which can also be controlled by the operator) to archive and retrieve data sets.

Though the existing MSS works, its dual operating system design has three major weaknesses. First, the development effort required to implement a working MSS with this design was considerable. Secondly, the extra communication required between UTS and MVS has converted UTS virtual disk I/O from an I/O bound process into a CPU bound process. Finally, since the MVS code is written in assembler language, the MSS software configuration consists of an unusually large number of lines of code.

Given these problems, the design of the MSS for EOC is based on two basic principles: use of a single operating system for efficiency and simplicity, and use of a high level language wherever possible.

The EOC MSS will be compatible with the other subsystems of the NPSN. Finally, since significant improvement over the current MSS can be achieved without changing the existing hardware, the EOC design for the MSS continues to use the Amdahl hardware configuration.

The design for the EOC MSS can be described as a series of upgrades to the hardware used by the current MSS , and a software re-implementation based on UTS/580. The target MSS software configuration for EOC is given in Figure 9 .

UTS/580 version 2.0

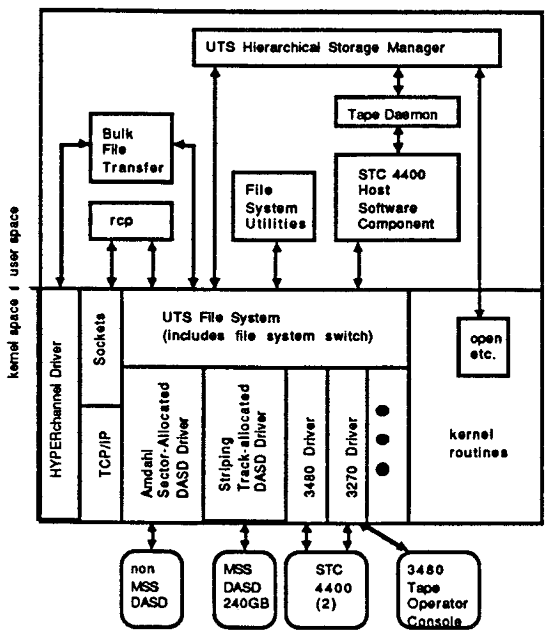

Figure 9. MSS EOC Software Configuration

Many of the design goals and requirements of the EOC MSS can be met by adding to and upgrading existing hardware. Improvements and additions will be made in several basic areas. The first area of improvement calls for an increase in available disk capacity. An initial upgrade in MSS disk space will increase existing disk storage from 120 GBytes to 240 GBytes. By the time of EOC, MSS disk capacity will be increased so that a minimum of 30-day residency of user files will be possible before archiving is required.

The EOC design for the MSS also includes a large improvement in CPU performance. This will be accomplished by reassigning both Amdahl machines to the MSS function. Both existing 5840 CPUs have been upgraded to 5860s for a 67% increase in processing speed, and then be joined to form a single 5880 multiprocessor. This change in the MSS configuration will require system software capable of exploiting a multiprocessor environment, although by 1988 both MVS and UTS are expected to provide this type of support. This basic change in the Amdahl configuration will provide 3.3 times the CPU power available to the existing MSS.

In conjunction with the changes to the Amdahl configuration, the HYPERchannel connection to the 'new' Amdahl 5880 will be upgraded to 4.5 MBytes/s, with corresponding changes to the HYPERchannel adapters. These new N-series adapters from Network Systems Corporation will be incorporated as soon as they are available.

In addition to increased disk capacity, the EOC MSS will incorporate an automated second level of storage below first-level disk. This second level of MSS storage at EOC will consist of at least two cartridge tape robots, the equivalent of Storage Technology's 4400. Theses devices will provide 2.4 Terrabytes of storage, accessible in an average of about 15 seconds. Since these units use dismountable media, they could also possibly satisfy the requirements for third-level storage.

The MSS EOC software configuration will be based on the UTS/580 operating system, release 2.0. Current information on the contents of this release indicates that it will include much of the software required for the EOC MSS . This release will :

1. Support File System Switch, enabling development of a striping track-allocated file system by NAS personnel

2. Support NFS, providing the distributed file system most likely to be adopted as the NPSN standard

3. Incorporate TCP/IP into the UTS design, with the probable result of a faster and more reliable implementation than the current Wollongong software

4. Support multiprocessor operation, enabling the MSS to exploit the entire 5880 processor

At least two additional software packages, running under UTS, will be incorporated into the EOC MSS. FileTek has developed a UNIX hierarchical file manager that is currently marketed in a system configuration based on a MicroVAX.

In addition, Storage Technology (STC) is currently working on the implementation of a UTS version of its Host Software Component, which helps control the STC cartridge tape robot. This software will also be implemented as an integral part of the EOC MSS.

To improve network file transfer rates, a bulk file transfer utility based on a HYPERchannel protocol will be essential to eliminating the TCP/IP, one-trunk HYPERchannel bottleneck. This file transfer utility will be available for MSS use in two stages. The first stage will support transfer of multiple files across multiple HYPERchannel adapter pairs, one or more files per pair. A later version of this utility will support use of multiple adapter pairs for a single file transfer.

4.3.4 Workstation Development Activities

The Workstation Subsystem (WKS) will include two kinds of workstations; the existing SGI 2500T Turbo and 3030 workstations (WK1), second generation graphics workstations (WK2)

WKS development activities are concerned with the integration of both kinds of workstation (WK1 & WK2). The necessary integration and development projects to support the graphics workstations in both the hardware and software areas are discussed below.

1. Acquisition, installation and initial operation of the most advanced scientific graphics workstations obtainable by competitive procurement.

2. A Proteon network hardware interface will be developed to evaluate the use of fiber optic token ring network technology in the NPSN environment.

3. A High Speed LAN interface will be developed to utilize the High Speed LAN being acquired as a part of the Data Network activities.

4. A graphics shell is being developed to act as an alternative to the 'C' and Bourne shells in the workstation environment.

5. A generalized standard interface is being developed to provide a common environment for software development and user interaction.

6. A front end interface and support tools will be provided for a Princeton University developed Navier Stokes Computer.

4.3.5 Advanced Graphics Effort Activities

The NAS graphics workstation effort is built around the principle that each scientist using NAS will require excellent numerical flow visualization tools. The principal tool to meet this requirements is the graphics workstations, WK1 and WK2, discussed in 4.4. The goal of the Advanced Graphics Effort (AGE) is to provide additional hardware and necessary software to exploit the capabilities of these workstations.

It is the intent of the NAS Program to support some research into the development of graphical techniques to identify new flow field structures. While basically graphical in nature, the research supported must be clearly applicable to CFD and must be of interest to the local ARC CFD community. The AGE will assist in funding hardware and salaries for such research

Hardware and software activities include:

1. Video signal digitization.

2. Incorporation of a high speed frame buffer,

3. Color printers & interfaces

4. RGB digital film recorders.

5. Advanced interfaces such as the use of "helmet" or "glove" interfaces to the WK's will be studied.

6. Remote graphics library and remote interactive particle tracing applications will be ported to WK2 and the Mini-Supercomputers.

7. Advanced interfaces including 3d space digitizers, helmets and gloves will be developed.

8. Visualization tools allowing the rendering experimental data and its direct comparison to CFD results.

9. Graphics animation system development providing stereo movies and ribbon traces.

10. The NAS Graphics Kernel benchmarks will be developed.

11. Porting standard applications including plot3d, gas, rip, and di3000.

4.3.6 Data Network Subsystem Development

The Data Network Subsystem (DNS) will provide for data communication service between NPSN subsystems. It will provide the connections and rates sufficient to support the required scientific computation and data management functions for NPSN users. The consistency and inter-operation of all communications software will be coordinated to assure that each processor is capable of communicating with the communications software of all other NPSN processors.

Conceptually, the DNS consists of both hardware and software. The DNS hardware includes physical media, media access and control hardware, physical adapters and buffers, host interface cards or equipment, and testing/monitoring/controlling hardware.

The DNS software includes microcode, firmware and software resident within DNS hardware, driver software within NPSN hosts, internet software including TCP/UDP/IP/ICMP, FTP and TELNET, Berkeley 4.3 BSD sockets, network library and r-commands, software to establish addresses and routes, specialized protocol software as required, and testing, monitoring and controlling software.

The DNS design for EOC is based on the following requirements for local data communications :

1. The DNS will provide a minimum data transfer rate of 2 megabits/sec. from memory to memory between any two NPSN processors.

2. The DNS will provide a data transfer rate of 8 megabits/sec. from memory to memory between an HSP and a new Workstation II.

3. The DNS will provide an aggregate data output rate of 100 megabits/sec. from an HSP.

4. The DNS will provide an aggregate data rate of 20 megabits/sec. between an HSP and the Mass Storage Subsystem.

5. The DNS will provide an aggregate data transfer rate of 32 megabits/sec. between all HSPs and the Long Haul Communications Subsystem.

6. The DNS will provide a point-to-point data transfer rate of 320 megabits/sec. between at least one HSP and the High Speed LAN or advanced graphics system.

7. Data within a packet or transmission unit will be delivered in the same sequence as transmitted.

8. No packet or transmission unit provided to the DNS will be delayed longer than 1 second.

9. The bit error rate will not exceed one error in every 10^9 bits.

10. The DNS media, adapters and interfaces will be reconfigurable to minimize loss of service due to any single component failure.

11. A DNS testing, monitoring and controlling system will be provided in at least the NAS Operators Room.

12. A performance monitoring system will be provided to accumulate data on traffic across the DNS.

For EOC, the DNS development effort will focus on four local area network technologies:

1. High Speed LAN
2. Proteon Ring
3. HYPERchannel
4. Ethernet

The High Speed LAN is a new, specialized system designed to provide the high rates of transfer required for graphics viewing or bulk file transfer. It is specified to provide a minimum of 320 megabits/sec.. It will connect the Cray 2, a graphics frame buffer, special graphics devices, and a small number of VME workstations. Additional connections to the HSP 2, MSS, APC, or VME-based systems are under study. It is expected that the High Speed LAN will use a specialized protocol.

The Proteon ring effort is a prototype activity whose objectives are to explore a high performance LAN, investigate fiber optic and token ring networking to gain experience for future application to FDDI-standard networks, and to examine the potential to concentrate multiple Ethernets or high speed lines onto a single LAN. If the prototype activities are successful, the Proteon network may be formally incorporated into the EOC configuration.

For EOC, the HYPERchannel network will be reconfigured to provide one main backbone network, as well as a second HSPs-to-MSS network. Both the new HSP 2 and the new mini-supercomputers will be attached to the HYPERchannel. For the HSPs and MSS the HYPERchannel adapters will be upgraded to new N-Series models offering higher performance. In addition, a dedicated unit serving as a HYPERchannel-to-Ethernet gateway will be installed.

For EOC, the main Ethernet network will be expanded to include connections to the Cray 2 and HSP 2, to the new workstations, and to the new mini-supercomputers. A second Ethernet line will be added to handle the high-speed remote user lines to the HSPs. A high speed bridge will be provided between the two Ethernet lines. Additional details on this reconfiguration are provided in the section that discusses LHC plans for EOC.

The Ethernet expansion and second line will be installed at the time of Ethernet installation for the Cray 2, in early 1988. The Ethernet expansion includes hardware for the Cray 2 and HSP 2 connections, a second Ethernet line for remote users, and bridge hardware for the connection between the two Ethernets.

The Ethernet expansion is expected to require minimal software additions. The Cray 2 Ethernet software will be provided by Cray Research.

A monitoring system will be attached to each of the above networks to collect network traffic and performance data. Additional software within each host will collect data on the inbound and outbound traffic of that host. LHC will also provide performance monitoring data.

Procedures to analyze this data will be prepared and performed to extract information needed for system analysis, evaluation of alternatives, load leveling, and emergency procedures.

Special hardware for each network will be required to accomplish this performance monitoring function. Performance monitors are currently available for the HYPERchannel and Ethernet networks. During the prototyping activities, methods for monitoring performance on the High Speed LAN and Proteon Ring networks will be developed.

Special software within the network monitoring hardware and the host computers will also be required. The specific requirements for this software depend on the characteristics to be monitored and the level of data collected. These detailed requirements have not yet been developed.

4.3.7 Long Haul Communications Subsystem EOC Design

The NAS Long Haul Communications Subsystem (LHCS) is responsible for providing high speed communications between the NPSN and its remote users.

The LHCS design for EOC is based on the following basic requirements:

1. With NAS Facility as the hub, NAS Long Haul Communications will provide the necessary support for simultaneous, direct (minimum hop) communication to remote user sites at the following data rates:

    a. One user operating at 6.312 megabits/sec. (T2)

    b. Five users operating at 1.544 megabits/sec. (T1)

    c. Twenty-five users operating at 56 kilobits/sec.

    d. Ten users operating at 9600 bps

2. NAS Long Haul Communications will support access to NAS via the following pathways:

    a. Milnet

    b. ARPAnet

    c. BARRnet (Bay Area Regional Research Network)

    d. NSFnet (National Science Foundation Network)

    e. Program Support Communications Network (PSCN)

        1). Computer Network Subsystem (CNS)

        2). NASA Packet Switched System (NPSS)

    f. Leased digital data services, including AT&T Accunet switched 56 service

    g. ARCLAN (Ames Research Center Local Area Network)

For those remote user sites with access to other networks that interconnect with NAS (such as Milnet and ARPAnet), this existing connectivity has been and will continue to be the method of choice for reaching the NAS computers. In addition, under an agreement with the National Science Foundation (NSF), university sites will be provided access to NAS via NSFnet.

Although remote communication to NAS emphasizes high-speed (56 kilobits/sec. and greater) workstation-to-supercomputer access, it is recognized that some low-speed connectivity will be necessary. For EOC, NAS Long Haul Communications will support 9.6 kilobits/sec. dial-up, dedicated service to selected University locations and other sites. The capability to support 10 of these connections will be available at EOC, although this does not necessarily imply that there will be ten 9.6 connections in actual use.

For those sites without existing connectivity and with requirements for high speed communications, such as industry sites, access to NAS will be provided via the remote network called NASnet. At the current time, the actual physical circuitry for remote communications connectivity to these users is implemented under the Program Support Communications (PSC) contract awarded to Boeing Computer Support Services (BCSS). The PSC contract is managed by Marshall Space Flight Center (MSFC).

The majority of the user sites receiving NASnet connections will be provided with switched 56 kilobits/sec. service. At the present time, 16 remote sites (industry locations and NASA centers) have this level of service. As implemented by BCSS, each industry location has a dedicated 56 kilobits/sec. AT&T circuit to the closest NASA center, where a channel on the PSCN backbone is then allocated for use (on demand) between that NASA center and ARC. Users at the NASA centers access NAS directly over the PSCN backbone (no AT&T interface). In addition to the switched 56 kilobits/sec. service, NASnet also includes dedicated 224 kilobits/sec. service to both Langley Research Center (LaRC) and Lewis Research Center (LeRC); the connection to LaRC is now being upgraded to 1.544 megabits/sec. (T1).

The configuration at each end of the NASnet circuits that allows computer-to-computer communication involves Ethernet bridging technology. NAS provides a Vitalink TransLAN Ethernet bridge to each location; attached to an Ethernet at the user site, the Vitalink equipment allows computers at the user site to communicate (via the AT&T and PSCN-provided circuitry) with computers at NAS via Vitalink equipment at NAS. Since all NAS computers employ the TCP/IP protocol suite, user-site computers must also run TCP/IP for communication to occur. Appropriate filters are implemented on each Vitalink bridge to limit access to only specific machines.

For future sites to be connected via NASnet, NAS will continue to provide a Vitalink TransLAN bridge to each site, as well as peripheral equipment such as a terminal, and an Ethernet transceiver. The LHC group will assist each site with installation. Each site will be brought on-line following a standardized procedure, in which the line is tested via DSU (Data Service Unit)-to-DSU loopback, followed by Vitalink-to-Vitalink connectivity, and finally Ethernet-to-Ethernet and host-to-host connectivity. Both the remote login (telnet) and file transfer (ftp) capabilities will also be tested. As each new site comes on-line, the circuit will be monitored by the LHC group for error rates and line availability, and all problems will be handled by the LHC group. Once the connection is considered stable, the circuit will be turned over to the NAS operations group (RNS). The expected time between the initial testing and turn-over to operations is approximately six weeks.

In addition to bringing new sites on-line, the NAS Long Haul Communications group will also arrange for upgrades in service to those locations with high activity over existing circuits. The Vitalink equipment at each remote site is capable of supporting rates up to 1.544 megabits/sec. and will not need to be replaced for these upgrades. The current schedule also calls for an upgrade to T2 (6.312 megabits/sec.) to LaRC, although Vitalink equipment to support this level of service is not yet available.

In order to detect communications problems, as well as to collect data on line usage for planning purposes, an automated system for monitoring the NASnet connections will be developed. The nucleus for this system exists at the present time, and consists of NAS-developed software running on one of the NAS VAXes. This software collects data traffic and error information from each of the Vitalink units every ten minutes, and stores this information in a file. Data are also being collected (manually) on a weekly basis from the PSCN switch, located in the ARC communications gateway building, as well as on each of the NAS computers using available accounting software. The goal for EOC is to combine these data collection capabilities in an automated system that will periodically generate reports on usage patterns and error trends. Currently the existing monitoring software runs on a VAX, but the VAXes will be phased out by EOC; a microVAX will serve as the Long Haul network monitoring machine for EOC.

One of the more significant findings of experiments conducted by the Long Haul Communications group is the problem of obtaining respectable throughput rates via satellite links using the TCP/IP protocol. Although sizable increases in throughput were obtained between two VAXes using identical TCP/IP implementations by manipulating buffer size and send/receive windows, the problem of widely differing TCP/IP implementations existing on a variety of different user machines makes this a non-general-purpose 'solution'. What will be required for efficient use of satellite capacity by NAS remote users is a 'tuned' protocol between the target NAS machines (both high speed processors and the Long Haul mini-supercomputer(s)), and a specific subset of machines at remote user locations. For the EOC timeframe, the subset of user machines will be limited to IRIS workstations running the same software as the NAS IRISes, and VAXes running either UNIX or VMS.

This task will require a thorough examination of the various TCP/IP implementations, as well as the investigation of other protocols, such as 'netblt' see ref. 3, in order to find an optimum protocol for use over satellite circuits.

The most recent version of TransLAN software available from Vitalink has the capability both to support parallel links to the same destination and to direct traffic over these parallel lines based on the type field in the Ethernet packet header. This capability will be prototyped to the Langley Research Center using the 1.544 megabits/sec. parallel terrestrial and satellite links.

For some time the local NPSN network support group has also been contemplating the implementation of TOS routing locally (some traffic is routed over HYPERchannel, other traffic over Ethernet), by specifying the TOS required in the IP packet header. These activities will be combined in order to provide improved service to both remote and local users. Coordination with the Vitalink Corporation will be required to modify the existing TransLAN software to utilize the IP rather than Ethernet packet header in designating TOS routing.

The results of the above prototype activities have immediate applicability for improving communications to remote NAS users and will be incorporated into the NPSN Long Haul environment as each task is successfully completed. However, the products of these activities will also be incorporated into an advanced prototype gateway.

The operational and conceptual paradigm for the advanced gateway prototype is its use in an interconnected network of networks ("internet"). NAS is currently connected to MILnet/ARPAnet, CNS, NPSS (NASA Packet Switched System, an X.25-based packet switched network), and the wideband services of NASnet. In the near future it will also have connectivity to such networks as NSFnet and the NASA Science Internet (NSI). With the capabilities provided by the PSCN, NAS is in an unique position to develop a prototype, advanced gateway. Such a gateway should be able to handle:

1. Switched as well as dedicated wideband circuits

2. Terrestrial (low delay) and satellite (high delay) service

3. Packet switched networks (e.g. MILnet, NPSS)

4. Interagency internets (e.g. NSFnet)

5. Expansion to bandwidths of T2 and greater

6. Multiple local network interfaces

7. Support for DoD-internet gateway protocols (ICMP, IP, EGP, etc.)

8. Type of service routing

9. Provide high-speed processing of packets (in the range of 10,000 packets per second)

The advanced gateway prototype activity will involve a number of steps, beginning with a formal definition of functional requirements. Once the requirements have been solidified, the design of the gateway itself will proceed; the design will be based on an existing IP gateway with 'standard' hardware interfaces and for which source code is available. The design process will be closely coordinated with one or more vendors currently producing viable gateway products.

The implementation process will include integration of the physical gateway with the dynamic T1 access methodology and the type of service routing capabilities described above. Performance of the gateway will then be measured and 'tuned'.

The ultimate goal of the gateway prototype is a specification for vendor implementation. Since a minimum of two gateways will be required to formally test the end system, at least one other NASA center will need to be involved (probably LaRC). Ideally, however, an additional center should also participate so that alternate routing and overload tests can be made.

4.3.8 Support Processing Subsystem Development

The current NAS Support Processors serve three classes of users. The first and largest class of users are Scientific researchers who have been granted use of NAS. The second group of users are NAS operational and administrative personnel, using the systems for support of a number of tasks including system administration, resource monitoring, and daily operational tasks. The third group of users are NAS development personnel who are doing software development, systems integration, testing, and performance monitoring. The major functions served by the Support Processors in the IOC are:

1. Terminal service - General support of terminals for full-screen text editing and login to other systems is provided.

2. Mail service - Electronic mail and messages.

3. Print service - Support for laser printers and other hard copy output devices.

4. Document preparation service - Including word processing and text editing capabilities.

5. Tape service - Including the reading, writing, cataloging and storage of magnetic tapes.

6. Administrative support - Development, administrative, and operational tasks necessary for the normal functioning of the Support Processors and the other production hosts in the NPSN.

The Support Processor functions are currently provided in the Initial Operating Configuration by three Vax 11/780's and one of two Amdahl 5860's. The second Amdahl 5860 supports the Mass Storage System function in NAS.

The IOC design approach will be modified by the time of the EOC for two major reasons. First, the demand for Mass Storage Resources will require both Amdahl processors to the exclusion of the current Support Processing Environment. The Amdahl 5860 system is currently shared between the Support Processor and Mass Storage function. Part of the plan for the Extended Operating Configuration is to provide a mass storage system with an aggregate transfer capacity of 25 megabits per second. Analysis of the current Mass Storage system, which provides an aggregate transfer rate of less than 2 megabits per second, indicate that the current configuration consumes nearly all of the available cycles on one of the CPU's. Although this analysis has also indicated that there are other bottlenecks that can be eliminated, it has indicated that the full capability of both current Amdahl CPU's will be needed to achieve the ambitious performance goal for the Extended Operating Configuration. Therefore, the Amdahl systems will not be available for the Support Processor function in the Extended Operating Configuration. Secondly, the growth of distributed interactive graphics applications for pre and post processing of CFD data presents an additional workload and performance requirement that cannot be met in the current environment. Examples of two distributed graphics applications are PLOT3D and RIP. PLOT3D is a locally developed graphics program that allows display and manipulation of grids and flowfields. RIP (which stands for Remote Interactive Particle Trace) allows NAS workstation users to interactively release particles and display their paths through a flowfield.

These distributed graphics applications have cooperating communicating processes on both the Cray 2 and the workstations. The processes running on the Cray-2 are responsible for most computation, and processes running on the workstations are responsible for the display of data and interaction with the user. Most current CFD calculations use grids with 500,000 to 1,000,000 grid points, and future problems will be larger. Distributed Cray-2 RIP processes use about 9 words of memory per grid point, while PLOT3D uses about 13 words of memory per grid point. Thus, at 8 bytes per word on the Cray-2, a PLOT3D process handling a 1,000,000 point grid would consume about 104 megabytes of memory, while a RIP process handling the same grid would require about 72 megabytes of memory.

While current graphics resource requirements are large the evolving requirement is even larger. There is a great interest in the CFD research community in being able to solve and analyze time-dependent problems. Solving a time-dependent problem generates a number of solution files. Each solution file represents the state of the model being studied at a particular point in time. A method of analysis of a sequence of these solutions through time is necessary. The Cray-2 is today capable of generating time-dependent solutions, but tools with the capacity to analyze the resulting data do not exist in NAS.

A time-dependent problem run today on the Cray-2 with 150,000 nodes, 8 variables per node, and 8 bytes per word on the Cray-2 would generate nearly 10 megabytes per solution. If every fifth solution were retained for 2000 iterations, the resulting 400 solution files would occupy about 4 gigabytes of disk space.

In order to examine this data and eventually make "movies" of the flowfield, the data must be available on a system where users can interactively examine and process it. The process of examining and processing this data shares many characteristics with the RIP and PLOT3D applications. Display of data and user interaction will take place on the workstations. Due to the volume of data that must be available readily, the solution files will have to be processed on another system that can communicate with the workstation. Users will initiate an action that changes the display of data on their workstations, and then typically spend some time thinking before requesting another action.

In summary, the current Support Processor applications are generally scalar in nature, are often input/output bound, and usually involve some sort of communication with other systems. The RIP and PLOT3D applications have proven to be a model for productive use of the NAS workstations. These applications consume a large amount of memory on the Cray-2, but do not use very much CPU time and are dependent on high performance network communication. The processing of results from time-dependent calculations, a major goal of NAS, will require that a very large amount of data be readily available from disk storage, and will share many of the characteristics of the RIP and PLOT3D applications.

A new set of systems are need in order to meet the Extended Operating Configuration Support Processor requirements. The new class of systems often called "mini-supercomputers" are best suited to addressing these requirements.

The current generation of graphics applications (in the form of RIP and PLOT3D) have already grown too large to be handled on NAS workstations. With increasing interest in time-dependent analysis, future graphics applications will require a much greater capability. The workstations are excellent in terms of display capability and user interaction, but become much more powerful when used in conjunction with other systems. Although the new generation of workstations that will be obtained for the Extended Operating Configuration will be much more powerful than the current ones, they will not have the storage or the speed to handle the analysis of data from time-dependent problems.

The power of the workstations has been extended in NAS with the development of distributed applications that connect the workstations directly to the Cray-2. As outlined in the previous section, these applications characteristically use a large amount of memory on the Cray-2, but are not heavily dependent on the fast CPU's that the Cray-2 offers. These applications are too large to be run on workstations, and the only NAS system that can process them at this point is the Cray-2. The mini-supercomputers offer a better solution for processing these applications than do supercomputers like the Cray-2 and future NAS high-speed processors. The reason for this lies in the characteristics of the applications as outlined in the previous section.

Thus, mini-supercomputers are an effective means of addressing current and future Support Processor applications. The role of systems that currently provide the Support Processor function will change as the new Support Processors are integrated into NAS.The overall requirements and performance levels for the mini-supercomputers are summarized in Table 3.

Table 3. Mini-supercomputer Requirements

Characteristic	Minimum	Target	Units
Peak Floating Point			
per CPU	5	10	Mflops
per System	10	20	Mflops
Sustained Performance	50%	50%	of above
Main Memory			
real	1	4	GBytes
virtual	2	4	GBytes
Disk Performance			
per device MBytes/sec	3	6	
aggragate MBytes/sec	3	24	
Configurable Disk	20	40	GBytes

Additional Features

Characteristic	Minimum	Target
Interfaces and Network Connections	9 track 1600/6250 bpi VME, Ethernet HYPERchannel	+HSX Pronet-80 3480 tape
Operating System	4.2/4.3 bsd and SUN NFS	+disk stripping
Networking Software	TCP/IP,ftp,telnet rcp,rsh on ethernet	+HYPERchannel
Languages	Fortran 77, C Pascal, Common Lisp	+Vectorizing C Fortran 8X

5.0 Applied Technology and Research Leading to NASII

In preparing for the future, evolution of the system past the design proposed for EOC is not dictated. Rather it is assumed that the system design might be radically different. The possible alternative architecture is designated NAS II and will allow NAS system designers to assume a blank sheet of paper in formulating the post-EOC system.

The design for NAS is formulated against two basic given requirements:

1. The demand for scientific computing capabilities will never be satisfied.

2. A fully productive scientific user interface does not exist and must be developed.

The current vision for NAS II currently contains four major elements:

1. Massive parallelism in both hardware and software.

2. A computational fluid dynamicists workbench containing graphical, numerical and symbolic computing capabilities.

3. Equality of access to both local and remote users of the NAS.

4. Distributed applications including distributed operating systems.

To clarify the vision for NAS II a applied research program and Advanced Development Laboratory is being established. The Laboratory will provide a testbed for investigating and evaluating future hardware and software from which enhancements to the NPSN beyond EOC can be developed. Areas of investigation include parallel processing architectures and algorithms, very high-speed local area networks, distributed operating systems, storage systems, and user interfaces employing workstations, visualization tools and expert systems technology.

6.0 Summary

The NAS Program has implemented the NPSN as a powerful supercomputing capability for the nation's aerospace research community. The Initial Operating Configuration was the first to achieve a supercomputing network of inhomogeneous hardware operating under UNIX as a single software environment. It pioneered the introduction of large memory supercomputers, exploited the power of scientific workstations, and provided high bandwidth access to researchers across the country. The NAS Program is now implementing the Extended Operating Configuration which will pioneer the introduction of the next generation supercomputer to the research community. State-or-the-art mass storage, networks , workstations and mini-supercomputers will be integrated into the EOC to achieve ready access, high utilization and maximum science benefit. Finally, the NAS Program is looking toward new systems beyond EOC and is defining a research program and Advanced Development Laboratory to meet future supercomputing challenges.

7.0 Acknowledgement

The authors would like to thank John Barton, Bill Kramer, Tom Lasinski, Judy McWilliams, Frank Preston, Karl Rowley and Dave Tweten for their valuable contributions in preparing this paper.

8.0 References

1. Peterson, Victor L.; and Ballhaus, William F., Jr: History of the Numerical Aerodynamic Simulation Program. Supercomputing in Aerospace, NASA CP 2454, March 1987, pp. 1-11.

2. Bailey, F. R.: NAS - Current Status and Future Plans. Supercomputing in Aerospace, NASA CP 2454, March 1987, pp. 13-21.

3. David A. Clark, "NETBLT: A bulk Data Transfer Protocol", RFC 998 March 1987.